

2. Introducing Differential Geometry

Gravity is geometry. To fully understand this statement, we will need more sophisticated tools and language to describe curved space and, ultimately, curved spacetime. This is the mathematical subject of *differential geometry* and will be introduced in this section and the next. Armed with these new tools, we will then return to the subject of gravity in Section 4.

Our discussion of differential geometry is not particularly rigorous. We will not prove many big theorems. Furthermore, a number of the statements that we make can be checked straightforwardly but we will often omit this. We will, however, be careful about building up the mathematical structure of curved spaces in the right logical order. As we proceed, we will come across a number of mathematical objects that can live on curved spaces. Many of these are familiar – like vectors, or differential operators – but we’ll see them appear in somewhat unfamiliar guises. The main purpose of this section is to understand what kind of objects can live on curved spaces, and the relationships between them. This will prove useful for both general relativity and other areas of physics.

Moreover, there is a wonderful rigidity to the language of differential geometry. It sometimes feels that any equation that you’re allowed to write down within this rigid structure is more likely than not to be true! This rigidity is going to be of enormous help when we return to discuss theories of gravity in Section 4.

2.1 Manifolds

The stage on which our story will play out is a mathematical object called a *manifold*. We will give a precise definition below, but for now you should think of a manifold as a curved, n -dimensional space. If you zoom in to any patch, the manifold looks like \mathbf{R}^n . But, viewed more globally, the manifold may have interesting curvature or topology.

To begin with, our manifold will have very little structure. For example, initially there will be no way to measure distances between points. But as we proceed, we will describe the various kinds of mathematical objects that can be associated to a manifold, and each one will allow us to do more and more things. It will be a surprisingly long time before we can measure distances between points! (Not until Section 3.)

You have met many manifolds in your education to date, even if you didn’t call them by name. Some simple examples in mathematics include Euclidean space \mathbf{R}^n , the sphere \mathbf{S}^n , and the torus $\mathbf{T}^n = \mathbf{S}^1 \times \dots \times \mathbf{S}^1$. Some simple examples in physics include the configuration space and phase space that we use in [classical mechanics](#) and

the state space of [thermodynamics](#). As we progress, we will see how familiar ideas in these subjects can be expressed in a more formal language. Ultimately our goal is to explain how spacetime is a manifold and to understand the structures that live on it.

2.1.1 Topological Spaces

Even before we get to a manifold, there is some work to do in order to define the underlying object. What follows is the mathematical equivalent of reading a biography about an interesting person and having to spend the first 20 pages wading through a description of what their grandparents did for a living. This backstory will not be particularly useful for our needs and we include it here only for completeness. We'll keep it down to one page.

Our backstory is called a *topological space*. Roughly speaking, this is a space in which each point can be viewed as living in a neighbourhood of other points, in a manner that allows us to define concepts such as continuity and convergence.

Definition: A *topological space* M is a set of points, endowed with a *topology* \mathcal{T} . This is a collection of open subsets $\{\mathcal{O}_\alpha \subset M\}$ which obey:

- i) Both the set M and the empty set \emptyset are open subsets: $M \in \mathcal{T}$ and $\emptyset \in \mathcal{T}$.
- ii) The intersection of a finite number of open sets is also an open set. So if $\mathcal{O}_1 \in \mathcal{T}$ and $\mathcal{O}_2 \in \mathcal{T}$ then $\mathcal{O}_1 \cap \mathcal{O}_2 \in \mathcal{T}$.
- iii) The union of any number (possibly infinite) of open sets is also an open set. So if $\mathcal{O}_\gamma \in \mathcal{T}$ then $\cup_\gamma \mathcal{O}_\gamma \in \mathcal{T}$.

Given a point $p \in M$, we say that $\mathcal{O} \in \mathcal{T}$ is a *neighbourhood* of p if $p \in \mathcal{O}$. This concept leads us to our final requirement: we require that, given any two distinct points, there is a neighbourhood which contains one but not the other. In other words, for any $p, q \in M$ with $p \neq q$, there exists $\mathcal{O}_1, \mathcal{O}_2 \in \mathcal{T}$ such that $p \in \mathcal{O}_1$ and $q \in \mathcal{O}_2$ and $\mathcal{O}_1 \cap \mathcal{O}_2 = \emptyset$. Topological spaces which obey this criterion are called *Hausdorff*. It is like a magic ward to protect us against bad things happening.

An example of a good Hausdorff space is the real line, $M = \mathbf{R}$, with \mathcal{T} consisting of all open intervals (a, b) , with $a < b \in \mathbf{R}$, and their unions. An example of a non-Hausdorff space is any M with $\mathcal{T} = \{M, \emptyset\}$.

Definition: One further definition (it won't be our last). A *homeomorphism* between topological spaces (M, \mathcal{T}) and $(\tilde{M}, \tilde{\mathcal{T}})$ is a map $f : M \rightarrow \tilde{M}$ which is

- i) Injective (or one-to-one): for $p \neq q$, $f(p) \neq f(q)$.
- ii) Surjective (or onto): $f(M) = \tilde{M}$, which means that for each $\tilde{p} \in \tilde{M}$ there exists a $p \in M$ such that $f(p) = \tilde{p}$.

Functions which are both injective and surjective are said to be bijective. This ensures that they have an inverse

- iii) Bicontinuous. This means that both the function and its inverse are continuous. To define a notion of continuity, we need to use the topology. We say that f is continuous if, for all $\tilde{\mathcal{O}} \in \tilde{\mathcal{T}}$, $f^{-1}(\tilde{\mathcal{O}}) \in \mathcal{T}$.

There's an animation of a donut morphing into a coffee mug and back that is often used to illustrate the idea of topology. If you want to be fancy, you can say that a donut is homeomorphic to a coffee mug.

2.1.2 Differentiable Manifolds

We now come to our main character: an n -dimensional manifold is a space which, locally, looks like \mathbf{R}^n . Globally, the manifold may be more interesting than \mathbf{R}^n , but the idea is that we can patch together these local descriptions to get an understanding for the entire space.

Definition: An n -dimensional *differentiable manifold* is a Hausdorff topological space M such that

- i) M is locally homeomorphic to \mathbf{R}^n . This means that for each $p \in M$, there is an open set \mathcal{O} such that $p \in \mathcal{O}$ and a homeomorphism $\phi : \mathcal{O} \rightarrow U$ with U an open subset of \mathbf{R}^n .
- ii) Take two open subsets \mathcal{O}_α and \mathcal{O}_β that overlap, so that $\mathcal{O}_\alpha \cap \mathcal{O}_\beta \neq \emptyset$. We require that the corresponding maps $\phi_\alpha : \mathcal{O}_\alpha \rightarrow U_\alpha$ and $\phi_\beta : \mathcal{O}_\beta \rightarrow U_\beta$ are *compatible*, meaning that the map $\phi_\beta \circ \phi_\alpha^{-1} : \phi_\alpha(\mathcal{O}_\alpha \cap \mathcal{O}_\beta) \rightarrow \phi_\beta(\mathcal{O}_\alpha \cap \mathcal{O}_\beta)$ is smooth (also known as infinitely differentiable or C^∞), as is its inverse. This is depicted in Figure 14.

The maps ϕ_α are called *charts* and the collection of charts is called an *atlas*. You should think of each chart as providing a coordinate system to label the region \mathcal{O}_α of M . The coordinate associated to $p \in \mathcal{O}_\alpha$ is

$$\phi_\alpha(p) = (x^1(p), \dots, x^n(p))$$

We write the coordinate in shorthand as simply $x^\mu(p)$, with $\mu = 1, \dots, n$. Note that we use a superscript μ rather than a subscript: this simple choice of notation will prove useful as we go along.

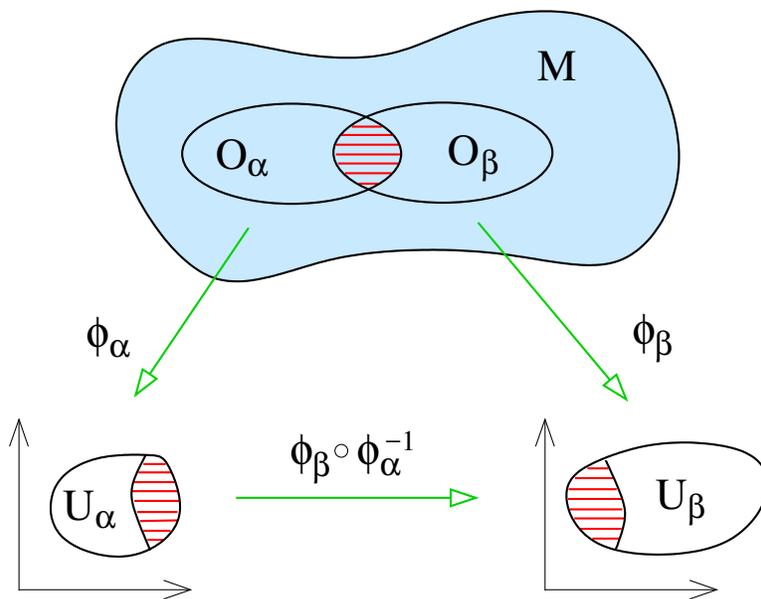


Figure 14: Charts on a manifold.

If a point p is a member of more than one subset \mathcal{O} then it may have a number of different coordinates associated to it. There's nothing to be nervous about here: it's entirely analogous to labelling a point using either Euclidean coordinate or polar coordinates.

The maps $\phi_\beta \circ \phi_\alpha^{-1}$ take us between different coordinate systems and are called *transition functions*. The compatibility condition is there to ensure that there is no inconsistency between these different coordinate systems.

Any manifold M admits many different atlases. In particular, nothing stops us from adding another chart to the atlas, provided that it is compatible with all the others. Two atlases are said to be compatible if every chart in one is compatible with every chart in the other. In this case, we say that the two atlases define the same *differentiable structure* on the manifold.

Examples

Here are a few simple examples of differentiable manifolds:

- \mathbf{R}^n : this looks locally like \mathbf{R}^n because it is \mathbf{R}^n . You only need a single chart with the usual Euclidean coordinates. Similarly, any open subset of \mathbf{R}^n is a manifold.

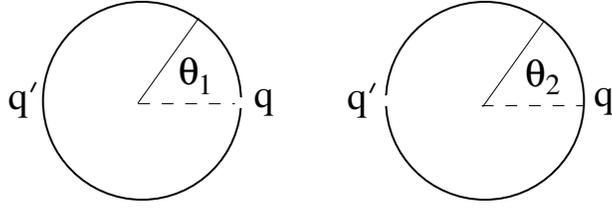


Figure 15: Two charts on a circle. The figures are subtly different! On the left, point q is removed and $\theta_1 \in (0, 2\pi)$. On the right, point q' is removed and $\theta_2 \in (-\pi, \pi)$.

- \mathbf{S}^1 : The circle can be defined as a curve in \mathbf{R}^2 with coordinates $(\cos \theta, \sin \theta)$. Until now in our physics careers, we've been perfectly happy taking $\theta \in [0, 2\pi)$ as the coordinate on \mathbf{S}^1 . But this coordinate does not meet our requirements to be a chart because it is not an open set. This causes problems if we want to differentiate functions at $\theta = 0$; to do so we need to take limits from both sides but there is no coordinate with θ a little less than zero.

To circumvent this, we need to use at least two charts to cover \mathbf{S}^1 . For example, we could pick out two antipodal points, say $q = (1, 0)$ and $q' = (-1, 0)$. We take the first chart to cover $\mathcal{O}_1 = \mathbf{S}^1 - \{q\}$ with the map $\phi_1 : \mathcal{O}_1 \rightarrow (0, 2\pi)$ defined by $\phi_1(p) = \theta$ as shown in the left-hand of Figure 15. We take the second chart to cover $\mathcal{O}_2 = \mathbf{S}^1 - \{q'\}$ with the map $\phi_2 : \mathcal{O}_2 \rightarrow (-\pi, \pi)$ defined by $\phi_2(p) = \theta'$ as shown in the right-hand figure.

The two charts overlap on the upper and lower semicircles. The transition function is given by

$$\theta' = \phi_2(\phi_1^{-1}(\theta)) = \begin{cases} \theta & \text{if } \theta \in (0, \pi) \\ \theta - 2\pi & \text{if } \theta \in (\pi, 2\pi) \end{cases}$$

The transition function isn't defined at $\theta = 0$, corresponding to the point q_1 , nor at $\theta = \pi$, corresponding to the point q_2 . Nonetheless, it is smooth on each of the two open intervals as required.

- \mathbf{S}^2 : It will be useful to think of the sphere as the surface $x^2 + y^2 + z^2 = 1$ embedded in Euclidean \mathbf{R}^3 . The familiar coordinates on the sphere \mathbf{S}^2 are those inherited from spherical polar coordinates of \mathbf{R}^3 , namely

$$x = \sin \theta \cos \phi, \quad y = \sin \theta \sin \phi, \quad z = \cos \theta \tag{2.1}$$

with $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. But as with the circle \mathbf{S}^1 described above, these are not open sets so will not do for our purpose. In fact, there are two distinct

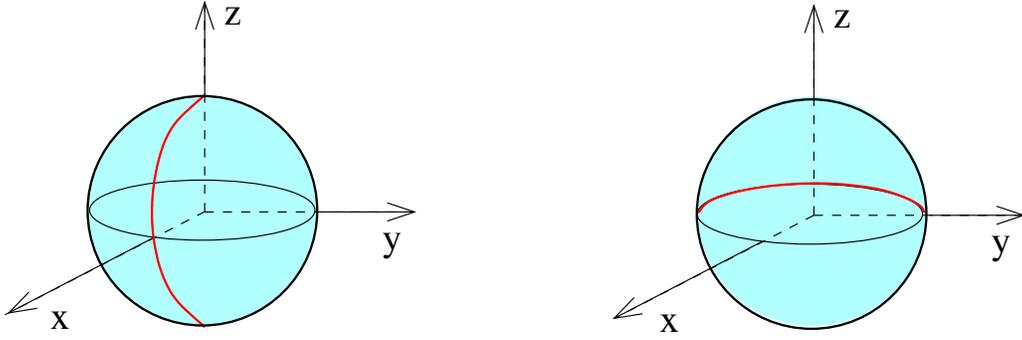


Figure 16: Two charts on a sphere. In the left-hand figure, we have removed the half-equator defined as $y = 0$ with $x > 0$, shown in red. In right-figure, we have removed the half-equator $z = 0$ with $x < 0$, again shown in red.

issues. If we focus on the equator at $\theta = \pi/2$, then the coordinate $\phi \in [0, 2\pi)$ parameterises a circle and suffers the same problem that we saw above. On top of this, at the north pole $\theta = 0$ and south pole $\theta = \pi$, the coordinate ϕ is not well defined, since the value of θ has already specified the point uniquely. This manifests itself on Earth by the fact that all time zones coincide at the North pole. It's one of the reasons people don't have business meetings there.

Once again, we can resolve these issues by introducing two charts covering different patches on \mathbf{S}^2 . The first chart applies to the sphere \mathbf{S}^2 with a line of longitude removed, defined by $y = 0$ and $x > 0$, as shown in Figure 16. (Think of this as the dateline.) This means that neither the north nor south pole are included in the open set \mathcal{O}_1 . On this open set, we define a map $\phi_1 : \mathcal{O}_1 \rightarrow \mathbf{R}^2$ using the coordinates (2.1), now with $\theta \in (0, \pi)$ and $\phi \in (0, 2\pi)$, so that we have a map to an open subset of \mathbf{R}^2 .

We then define a second chart on a different open set \mathcal{O}_2 , defined by \mathbf{S}^2 , with the line $z = 0$ and $x < 0$ removed. Here we define the map $\phi_2 : \mathcal{O}_2 \rightarrow \mathbf{R}^2$ using the coordinates

$$x = -\sin \theta' \cos \phi' \quad , \quad y = \cos \theta' \quad , \quad z = \sin \theta' \sin \phi'$$

with $\theta' \in (0, \pi)$ and $\phi \in (0, 2\pi)$. Again this is a map to an open subset of \mathbf{R}^2 . We have $\mathcal{O}_1 \cup \mathcal{O}_2 = \mathbf{S}^2$ while, on the overlap $\mathcal{O}_1 \cap \mathcal{O}_2$, the transition functions $\phi_1 \circ \phi_2^{-1}$ and $\phi_2 \circ \phi_1^{-1}$ are smooth. (We haven't written these functions down explicitly, but it's clear that they are built from cos and sin functions acting on domains where their inverses exist.)

Note that for both \mathbf{S}^1 and \mathbf{S}^2 examples above, we made use of the fact that they can be viewed as embedded in a higher dimensional \mathbf{R}^{n+1} to construct the charts. However, this isn't necessary. The definition of a manifold makes no mention of a higher dimensional embedding and these manifolds should be viewed as having an existence independent of any embedding.

As you can see, there is a level of pedantry involved in describing these charts. (Mathematicians prefer the word “rigour”.) The need to deal with multiple charts arises only when we have manifolds of non-trivial topology; the manifolds \mathbf{S}^1 and \mathbf{S}^2 that we met above are particularly simple examples. When we come to discuss general relativity, we will care a lot about changing coordinates, and the limitations of certain coordinate systems, but our manifolds will turn out to be simple enough that, for all practical purposes, we can always find a single set of coordinates that tells us what we need to know. However, as we progress in physics, and topology becomes more important, so too does the idea of different charts. Perhaps the first place in physics where overlapping charts become an integral part of the discussion is the construction of a magnetic monopole. (See the lectures on [Gauge Theory](#).)

2.1.3 Maps Between Manifolds

The advantage of locally mapping a manifold to \mathbf{R}^n is that we can now import our knowledge of how to do maths on \mathbf{R}^n . For example, we know how to differentiate functions on \mathbf{R}^n , and what it means for functions to be smooth. This now translates directly into properties of functions defined over the manifold.

We say that a function $f : M \rightarrow \mathbf{R}$ is smooth, if the map $f \circ \phi^{-1} : U \rightarrow \mathbf{R}$ is smooth for all charts ϕ .

Similarly, we say that a map $f : M \rightarrow N$ between two manifolds M and N (which may have different dimensions) is smooth if the map $\psi \circ f \circ \phi^{-1} : U \rightarrow V$ is smooth for all charts $\phi : M \rightarrow U \subset \mathbf{R}^{\dim(M)}$ and $\psi : N \rightarrow V \subset \mathbf{R}^{\dim(N)}$

A *diffeomorphism* is defined to be a smooth homeomorphism $f : M \rightarrow N$. In other words it is an invertible, smooth map between manifolds M and N that has a smooth inverse. If such a diffeomorphism exists then the manifolds M and N are said to be *diffeomorphic*. The existence of an inverse means M and N necessarily have the same dimension.

Manifolds which are homeomorphic can be continuously deformed into each other. But diffeomorphism is stronger: it requires that the map and its inverse are smooth. This gives rise to some curiosities. For example, it turns out that the sphere \mathbf{S}^7 can

be covered by a number of different, incompatible atlases. The resulting manifolds are homeomorphic but not diffeomorphic. These are referred to as exotic spheres. Similarly, Euclidean space \mathbf{R}^n has a unique differentiable structure, except for \mathbf{R}^4 where there are an infinite number of inequivalent structures. I know of only one application of exotic spheres to physics (a subtle global gravitational anomaly in superstring theory) and I know of no applications of the exotic differential structure on \mathbf{R}^4 . Certainly these will not play any role in these lectures.

2.2 Tangent Spaces

Our next task is to understand how to do calculus on manifolds. We start here with differentiation; it will take us a while longer to get to integration, which we will finally meet in Section 2.4.4.

Consider a function $f : M \rightarrow \mathbf{R}$. To differentiate the function at some point p , we introduce a chart $\phi = (x^1, \dots, x^n)$ in a neighbourhood of p . We can then construct the map $f \circ \phi^{-1} : U \rightarrow \mathbf{R}$ with $U \subset \mathbf{R}^n$. But we know how to differentiate functions on \mathbf{R}^n and this gives us a way to differentiate functions on M , namely

$$\left. \frac{\partial f}{\partial x^\mu} \right|_p := \left. \frac{\partial (f \circ \phi^{-1})}{\partial x^\mu} \right|_{\phi(p)} \quad (2.2)$$

Clearly this depends on the choice of chart ϕ and coordinates x^μ . We would like to give a coordinate independent definition of differentiation, and then understand what happens when we choose to describe this object using different coordinates.

2.2.1 Tangent Vectors

We will consider smooth functions over a manifold M . We denote the set of all smooth functions as $C^\infty(M)$.

Definition: A *tangent vector* X_p is an object that differentiates functions at a point $p \in M$. Specifically, $X_p : C^\infty(M) \rightarrow \mathbf{R}$ satisfying

- i) Linearity: $X_p(f + g) = X_p(f) + X_p(g)$ for all $f, g \in C^\infty(M)$.
- ii) $X_p(f) = 0$ when f is the constant function.
- iii) Leibnizarity: $X_p(fg) = f(p)X_p(g) + X_p(f)g(p)$ for all $f, g \in C^\infty(M)$. This, of course, is the product rule.

Note that ii) and iii) combine to tell us that $X_p(af) = aX_p(f)$ for $a \in \mathbf{R}$.

This definition is one of the early surprises in differential geometry. The surprise is really in the name “tangent vector”. We know what vectors are from undergraduate physics, and we know what differential operators are. But we’re not used to equating the two. Before we move on, it might be useful to think about how this definition fits with other notions of vectors that we’ve met before.

The first time we meet a vector in physics is usually in the context of Newtonian mechanics, where we describe the position of a particle as a vector \mathbf{x} in \mathbf{R}^3 . This concept of a vector is special to flat space and does not generalise to other manifolds. For example, a line connecting two points on a sphere is not a vector and, in general, there is no way to think of a point $p \in M$ as a vector. So we should simply forget that points in \mathbf{R}^3 can be thought of as vectors.

The next type of vector is the velocity of a particle, $\mathbf{v} = \dot{\mathbf{x}}$. This is more pertinent. It clearly involves differentiation and, moreover, is tangent to the curve traced out by the particle. As we will see below, velocities of particles are indeed examples of tangent vectors in differential geometry. More generally, tangent vectors tell us how things change in a given direction. They do this by differentiating.

It is simple to check that the object

$$\partial_\mu \Big|_p := \frac{\partial}{\partial x^\mu} \Big|_p$$

which acts on functions as shown in (2.2) obeys all the requirements of a tangent vector.

Note that the index μ is now a subscript, rather than superscript that we used for the coordinates x^μ . (On the right-hand-side, the superscript in $\partial/\partial x^\mu$ is in the denominator and counts as a subscript.) We will adopt the summation convention, where repeated indices are summed. But, as we will see, the placement of indices up or down will tell us something and all sums will necessarily have one index up and one index down. This is a convention that we met already in [Special Relativity](#) where the up/downness of the index changes minus signs. Here it has a more important role that we will see as we go on: the placement of the index tells us what kind of mathematical space the object lives in. For now, you should be aware that any equation with two repeated indices that are both up or both down is necessarily wrong, just as any equation with three or more repeated indices is wrong.

Theorem: The set of all tangent vectors at point p forms an n -dimensional vector space. We call this the *tangent space* $T_p(M)$. The tangent vectors $\partial_\mu \Big|_p$ provide a basis

for $T_p(M)$. This means that we can write any tangent vector as

$$X_p = X^\mu \partial_\mu \Big|_p$$

with $X^\mu = X_p(x^\mu)$ the components of the tangent vector in this basis.

Proof: Much of the proof is just getting straight what objects live in what spaces. Indeed, getting this straight is a large part of the subject of differential geometry. To start, we need a small lemma. We define the function $F = f \circ \phi^{-1} : U \rightarrow \mathbf{R}$, with $\phi = (x^1, \dots, x^n)$ a chart on a neighbourhood of p . Then, in some (perhaps smaller) neighbourhood of p we can always write the function F as

$$F(x) = F(x^\mu(p)) + (x^\mu - x^\mu(p))F_\mu(x) \quad (2.3)$$

where we have introduced n new functions $F_\mu(x)$ and used the summation convention in the final term. If the function F has a Taylor expansion then we can trivially write it in the form (2.3) by repackaging all the terms that are quadratic and higher into the $F_\mu(x)$ functions, keeping a linear term out front. But in fact there's no need to assume the existence of a Taylor expansion. One way to see this is to note that for any function $G(t)$ we trivially have $G(1) = G(0) + \int_0^1 dt G'(t)$. But now apply this formula to the function $G(t) = F(tx)$ for some fixed x . This gives $F(x) = F(0) + x \int_0^1 dt F'(xt)$ which is precisely (2.3) for a function of a single variable expanded about the origin. The same method holds more generally.

Given (2.3), we act with ∂_μ on both sides, and then evaluate at $x^\mu = x^\mu(p)$. This tells us that the functions F_μ must satisfy

$$\frac{\partial F}{\partial x^\mu} \Big|_{x(p)} = F_\mu(x(p)) \quad (2.4)$$

We can translate this into a similar expression for f itself. We define n functions on M by $f_\mu = F_\mu \circ \phi$. Then, for any $q \in M$ in the appropriate neighbourhood of p , (2.3) becomes

$$f \circ \phi^{-1}(x^\mu(q)) = f \circ \phi^{-1}(x^\mu(p)) + (x^\mu(q) - x^\mu(p)) [f_\mu \circ \phi^{-1}(x^\mu(q))]$$

But $\phi^{-1}(x^\mu(q)) = q$. So we find that, in the neighbourhood of p , it is always possible to write a function f as

$$f(q) = f(p) + (x^\mu(q) - x^\mu(p))f_\mu(q)$$

for some $f_\mu(q)$. Note that, evaluated at $q = p$, we have

$$f_\mu(p) = F_\mu \circ \phi(p) = F_\mu(x(p)) = \frac{\partial F}{\partial x^\mu} \Big|_{x(p)} = \frac{\partial f}{\partial x^\mu} \Big|_p$$

where in the last equality we used (2.2) and in the penultimate equality we used (2.4).

Now we can turn to the tangent vector X_p . This acts on the function f to give

$$X_p(f) = X_p\left(f(p) + (x^\mu - x^\mu(p))f_\mu\right)$$

where we've dropped the arbitrary argument q in $f(q)$, $x^\mu(q)$ and $f_\mu(q)$; these are the functions on which the tangent vector is acting. Using linearity and Leibnizarity, we have

$$X_p(f) = X_p\left(f(p)\right) + X_p\left((x^\mu - x^\mu(p))\right)f_\mu(p) + (x^\mu(p) - x^\mu(p))X_p\left(f_\mu\right)$$

The first term vanishes because $f(p)$ is just a constant and all tangent vectors are vanishing when acting on a constant. The final term vanishes as well because the Leibniz rule tells us to evaluate the function $(x^\mu - x^\mu(p))$ at p . Finally, by linearity, the middle term includes a $X_p(x^\mu(p))$ term which vanishes because $x^\mu(p)$ is just a constant. We're left with

$$X_p(f) = X_p(x^\mu) \left. \frac{\partial f}{\partial x^\mu} \right|_p$$

This means that the tangent vector X_p can be written as

$$X_p = X^\mu \left. \frac{\partial}{\partial x^\mu} \right|_p$$

with $X^\mu = X_p(x^\mu)$ as promised. To finish, we just need to show that $\left. \partial_\mu \right|_p$ provide a basis for $T_p(M)$. From above, they span the space. To check linear independence, suppose that we have vector $\alpha = \alpha_\mu \left. \partial_\mu \right|_p = 0$. Then acting on $f = x^\nu$, this gives $\alpha(x^\nu) = \alpha_\mu (\partial_\mu x^\nu) \Big|_p = \alpha_\nu = 0$. This concludes our proof. \square

Changing Coordinates

We have an ambivalent relationship with coordinates. We can't calculate anything without them, but we don't want to rely on them. The compromise we will come to is to consistently check that nothing physical depends on our choice of coordinates.

The key idea is that a given tangent vector X_p exists independent of the choice of coordinate. However, the chosen basis $\{\left. \partial_\mu \right|_p\}$ clearly depends on our choice of coordinates: to define it we had to first introduce a chart ϕ and coordinates x^μ . A basis defined in this way is called, quite reasonably, a *coordinate basis*. At times we will work with other bases, $\{e_\mu\}$ which are not defined in this way. Unsurprisingly, these are referred to as *non-coordinate bases*. A particularly useful example of a non-coordinate basis, known as *vielbeins*, will be introduced in Section 3.4.2.

Suppose that we picked a different chart $\tilde{\phi}$, with coordinates \tilde{x}^μ in the neighbourhood of p . We then have two different bases, and can express the tangent vector X_p in terms of either,

$$X_p = X^\mu \left. \frac{\partial}{\partial x^\mu} \right|_p = \tilde{X}^\mu \left. \frac{\partial}{\partial \tilde{x}^\mu} \right|_p$$

The vector is the same, but the components of the vector change: they are X^μ in the first set of coordinates, and \tilde{X}^μ in the second. It is straightforward to determine the relationship between X^μ and \tilde{X}^μ . To see this, we look at how the tangent vector X_p acts on a function f ,

$$X_p(f) = X^\mu \left. \frac{\partial f}{\partial x^\mu} \right|_p = X^\mu \left. \frac{\partial \tilde{x}^\nu}{\partial x^\mu} \right|_{\phi(p)} \left. \frac{\partial f}{\partial \tilde{x}^\nu} \right|_p$$

where we've used the chain rule. (Actually, we've been a little quick here. You can be more careful by introducing the functions $F = f \circ \phi^{-1}$ and $\tilde{F} = f \circ \tilde{\phi}^{-1}$ and using (2.2) to write $\partial f / \partial x^\mu = \partial F(\tilde{x}(x)) / \partial x^\mu$. The end result is the same. We will be similarly sloppy in the same way as we proceed, often conflating f and F .) You can read this equation in one of two different ways. First, we can view this as a change in the basis vectors: they are related as

$$\left. \frac{\partial}{\partial x^\mu} \right|_p = \left. \frac{\partial \tilde{x}^\nu}{\partial x^\mu} \right|_{\phi(p)} \left. \frac{\partial}{\partial \tilde{x}^\nu} \right|_p \quad (2.5)$$

Alternatively, we can view this as a change in the components of the vector, which transform as

$$\tilde{X}^\nu = X^\mu \left. \frac{\partial \tilde{x}^\nu}{\partial x^\mu} \right|_{\phi(p)} \quad (2.6)$$

Components of vectors that transform this way are sometimes said to be *contravariant*. I've always found this to be annoying terminology, in large part because I can never remember it. A more important point is that the form of (2.6) is essentially fixed once you remember that the index on X^μ sits up rather than down.

What Is It Tangent To?

So far, we haven't really explained where the name "tangent vector" comes from. Consider a smooth curve in M that passes through the point p . This is a map $\sigma : I \rightarrow M$, with I an open interval $I \subset \mathbf{R}$. We will parameterise the curve as $\sigma(t)$ such that $\sigma(0) = p \in M$.

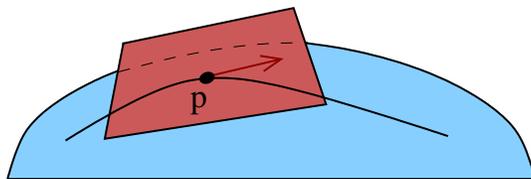


Figure 17: The tangent space at a point p .

With a given chart, this curve becomes $\phi \circ \sigma : \mathbf{R} \mapsto \mathbf{R}^n$, parameterised by $x^\mu(t)$. Before we learned any differential geometry, we would say that the tangent vector to the curve at $t = 0$ is

$$X^\mu = \left. \frac{dx^\mu(t)}{dt} \right|_{t=0}$$

But we can take these to be the components of the tangent vector X_p , which we define as

$$X_p = \left. \frac{dx^\mu(t)}{dt} \right|_{t=0} \left. \frac{\partial}{\partial x^\mu} \right|_p$$

Our tangent vector now acts on functions $f \in C^\infty(M)$. It is telling us how fast any function f changes as we move along the curve.

Any tangent vector X_p can be written in this form. This gives meaning to the term “tangent space” for $T_p(M)$. It is, literally, the space of all possible tangents to curves passing through the point p . For example, a two dimensional manifold, embedded in \mathbf{R}^3 is shown in Figure 17. At each point p , we can identify a vector space which is the tangent plane: this is $T_p(M)$.

As an aside, note that the mathematical definition of a tangent space makes no reference to embedding the manifold in some higher dimensional space. The tangent space is an object intrinsic to the manifold itself. (This is in contrast to the picture where it was unfortunately necessary to think about the manifold as embedded in \mathbf{R}^3 .)

The tangent spaces $T_p(M)$ and $T_q(M)$ at different points $p \neq q$ are different. There’s no sense in which we can add vectors from one to vectors from the other. In fact, at this stage there’s no way to even compare vectors in $T_p(M)$ to vectors in $T_q(M)$. They are simply different spaces. As we proceed, we will make some effort to figure ways to get around this.

2.2.2 Vector Fields

So far we have only defined tangent vectors at a point p . It is useful to consider an object in which there is a choice of tangent vector for every point $p \in M$. In physics, we call objects that vary over space *fields*.

A *vector field* X is defined to be a smooth assignment of a tangent vector X_p to each point $p \in M$. This means that if you feed a function to a vector field, then it spits back another function, which is the differentiation of the first. In symbols, a vector field is therefore a map $X : C^\infty(M) \rightarrow C^\infty(M)$. The function $X(f)$ is defined by

$$(X(f))(p) = X_p(f)$$

The space of all vector fields on M is denoted $\mathfrak{X}(M)$.

Given a coordinate basis, we can expand any vector field as

$$X = X^\mu \frac{\partial}{\partial x^\mu} \tag{2.7}$$

where the X^μ are now smooth functions on M .

Strictly speaking, the expression (2.7) only defines a vector field on the open set $\mathcal{O} \subset M$ covered by the chart, rather than the whole manifold. We may have to patch this together with other charts to cover all of M .

The Commutator

Given two vector fields $X, Y \in \mathfrak{X}(M)$, we can't multiply them together to get a new vector field. Roughly speaking, this is because the product XY is a second order differential operator rather than a first order operator. This reveals itself in a failure of Leibnizarity for the object XY ,

$$XY(fg) = X(fY(g) + Y(f)g) = X(f)Y(g) + fXY(g) + gXY(f) + X(g)Y(f)$$

This is not the same as $fXY(g) + gXY(f)$ that Leibniz requires.

However, we can build a new vector field by taking the *commutator* $[X, Y]$, which acts on functions f as

$$[X, Y](f) = X(Y(f)) - Y(X(f))$$

This is also known as the *Lie bracket*. Evaluated in a coordinate basis, the commutator is given by

$$\begin{aligned} [X, Y](f) &= X^\mu \frac{\partial}{\partial x^\mu} \left(Y^\nu \frac{\partial f}{\partial x^\nu} \right) - Y^\mu \frac{\partial}{\partial x^\mu} \left(X^\nu \frac{\partial f}{\partial x^\nu} \right) \\ &= \left(X^\mu \frac{\partial Y^\nu}{\partial x^\mu} - Y^\mu \frac{\partial X^\nu}{\partial x^\mu} \right) \frac{\partial f}{\partial x^\nu} \end{aligned}$$

This holds for all $f \in C^\infty(M)$, so we're at liberty to write

$$[X, Y] = \left(X^\mu \frac{\partial Y^\nu}{\partial x^\mu} - Y^\mu \frac{\partial X^\nu}{\partial x^\mu} \right) \frac{\partial}{\partial x^\nu} \quad (2.8)$$

It is not difficult to check that the commutator obeys the Jacobi identity

$$[X, [Y, Z]] + [Y, [Z, X]] + [Z, [X, Y]] = 0$$

This ensures that the set of all vector fields on a manifold M has the mathematical structure of a Lie algebra.

2.2.3 Integral Curves

There is a slightly different way of thinking about vector fields on a manifold. A *flow* on M is a one-parameter family of diffeomorphisms $\sigma_t : M \rightarrow M$ labelled by $t \in \mathbf{R}$. These maps have the properties that $\sigma_{t=0}$ is the identity map, and $\sigma_s \circ \sigma_t = \sigma_{s+t}$. These two requirements ensure that $\sigma_{-t} = \sigma_t^{-1}$. Such a flow gives rise to streamlines on the manifold. We will further require that these streamlines are smooth.

We can then define a vector field by taking the tangent to the streamlines at each point. In a given coordinate system, the components of the vector field are

$$X^\mu(x^\mu(t)) = \frac{dx^\mu(t)}{dt} \quad (2.9)$$

where I've abused notation a little and written $x^\mu(t)$ rather than the more accurate but cumbersome $x^\mu(\sigma_t)$. This will become a habit, with the coordinates x^μ often used to refer to the point $p \in M$.

A flow gives rise to a vector field. Alternatively, given a vector field $X^\mu(x)$, we can integrate the differential equation (2.9), subject to an initial condition $x^\mu(0) = x_{\text{initial}}^\mu$ to generate streamlines which start at x_{initial}^μ . These streamlines are called *integral curves*, generated by X .

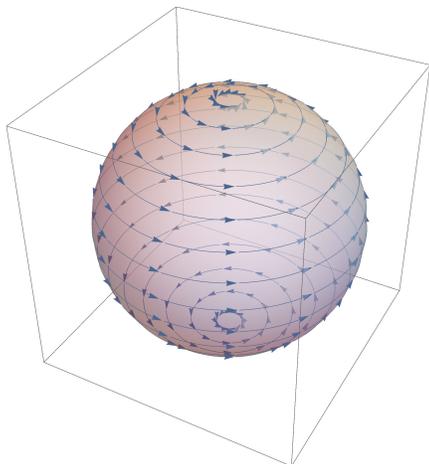


Figure 18: Flows on a sphere.

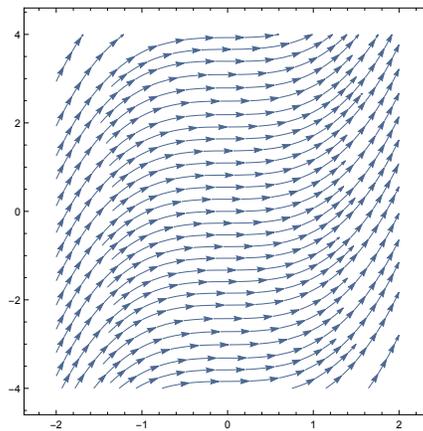


Figure 19: Flows in the plane.

In what follows, we will only need the infinitesimal flow generated by X . This is simply

$$x^\mu(t) = x^\mu(0) + tX^\mu(x) + \mathcal{O}(t^2) \quad (2.10)$$

Indeed, differentiating this obeys (2.9) to leading order in t .

(An aside: Given a vector field X , it may not be possible to integrate (2.9) to generate a flow defined for all $t \in \mathbf{R}$. For example, consider $M = \mathbf{R}$ with the vector field $X = x^2$. The equation $dx/dt = x^2$, subject to the initial condition $x(0) = a$, has the unique solution $x(t) = a/(1 - at)$ which diverges at $t = 1/a$. Vector fields which generate a flow for all $t \in \mathbf{R}$ are called *complete*. It turns out that all vector fields on a manifold M are complete if M is *compact*. Roughly speaking, “compact” means that M doesn’t “stretch to infinity”. More precisely, a topological space M is compact if, for any family of open sets covering M there always exists a finite sub-family which also cover M . So \mathbf{R} is not compact because the family of sets $\{(-n, n), n \in \mathbf{Z}^+\}$ covers \mathbf{R} but has no finite sub-family. Similarly, \mathbf{R}^n is non-compact. However, \mathbf{S}^n and \mathbf{T}^n are compact manifolds.)

We can look at some examples.

- Consider the sphere \mathbf{S}^2 in polar coordinates with the vector field $X = \partial_\phi$. The integral curves solve the equation (2.9), which are

$$\frac{d\phi}{dt} = 1 \quad \text{and} \quad \frac{d\theta}{dt} = 0$$

This has the solution $\theta = \theta_0$ and $\phi = \phi_0 + t$. The associated one-parameter diffeomorphism is $\sigma_t : (\theta, \phi) \rightarrow (\theta, \phi + t)$, and the flow lines are simply lines of constant latitude on the sphere and are shown in the left-hand figure.

- Alternatively, consider the vector field on \mathbf{R}^2 with Cartesian components $X^\mu = (1, x^2)$. The equation for the integral curves is now

$$\frac{dx}{dt} = 1 \quad \text{and} \quad \frac{dy}{dt} = x^2$$

which has the solution $x(t) = x_0 + t$ and $y(t) = y_0 + \frac{1}{3}(x_0 + t)^3$. The associated flow lines are shown in the right-hand figure.

2.2.4 The Lie Derivative

So far we have learned how to differentiate a function. This requires us to introduce a vector field X , and the new function $X(f)$ can be viewed as the derivative of f in the direction of X .

Next we ask: is it possible to differentiate a vector field? Specifically, suppose that we have a second vector field Y . How can we differentiate this in the direction of X to get a new vector field? As we've seen, we can't just write down XY because this doesn't define a new vector field.

To proceed, we should think more carefully about what differentiation means. For a function $f(x)$ on \mathbf{R} , we compare the values of the function at nearby points, and see what happens as those points approach each other

$$\frac{df}{dx} = \lim_{t \rightarrow 0} \frac{f(x+t) - f(x)}{t}$$

Similarly, to differentiate a vector field, we need to subtract the tangent vector $Y_p \in T_p(M)$ from the tangent vector at some nearby point $Y_q \in T_q(M)$, and then see what happens in the limit $q \rightarrow p$. But that's problematic because, as we stressed above, the vector spaces $T_p(M)$ and $T_q(M)$ are different, and it makes no sense to subtract vectors in one from vectors in the other. To make progress, we're going to have to find a way to do this. Fortunately, there is a way.

Push-Forward and Pull-Back

Suppose that we have a map $\varphi : M \rightarrow N$ between two manifolds M and N which we will take to be a diffeomorphism. This allows us to import various structures on one manifold to the other.

For example, if we have a function on $f : N \rightarrow \mathbf{R}$, then we can construct a new function that we denote $(\varphi^* f) : M \rightarrow \mathbf{R}$,

$$(\varphi^* f)(p) = f(\varphi(p))$$

Using the map in this way, to drag objects originally defined on N onto M is called the *pull-back*. If we introduce coordinates x^μ on M and y^α on N , then the map $\varphi(x) = y^\alpha(x)$, and we can write

$$(\varphi^* f)(x) = f(y(x))$$

Some objects more naturally go the other way. For example, given a vector field Y on M , we can define a new vector field $(\varphi_* Y)$ on N . If we are given a function $f : N \rightarrow \mathbf{R}$, then the vector field $(\varphi_* Y)$ on N acts as

$$(\varphi_* Y)(f) = Y(\varphi^* f)$$

where I've been a little sloppy in the notation here since the left-hand side is a function on N and the right-hand side a function on M . The equality above holds when evaluated at the appropriate points: $[(\varphi_* Y)(f)](\varphi(p)) = [Y(\varphi^* f)](p)$. Using the map to push objects on M onto N is called the *push-forward*.

If $Y = Y^\mu \partial / \partial x^\mu$ is the vector field on M , we can write the induced vector field on N as

$$(\varphi_* Y)(f) = Y^\mu \frac{\partial f(y(x))}{\partial x^\mu} = Y^\mu \frac{\partial y^\alpha}{\partial x^\mu} \frac{\partial f(y)}{\partial y^\alpha}$$

Written in components, $(\varphi_* Y) = (\varphi_* Y)^\alpha \partial / \partial y^\alpha$, we then have

$$(\varphi_* Y)^\alpha = Y^\mu \frac{\partial y^\alpha}{\partial x^\mu} \tag{2.11}$$

Given the way that the indices are contracted, this is more or less the only thing we could write down.

We'll see other examples of these induced maps later in the lectures. The push-forward is always denoted as φ_* and goes in the same way as the original map. The pull-back is always denoted as φ^* and goes in the opposite direction to the original map. Importantly, if our map $\varphi : M \rightarrow N$ is a diffeomorphism, then we also have $\varphi^{-1} : N \rightarrow M$, so we can transport any object from M to N and back again with impunity.

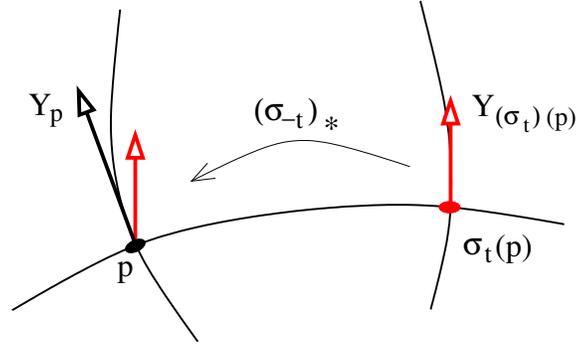


Figure 20: To construct the Lie derivative, we use the push-forward $(\sigma_{-t})_*$ to map the vector $Y_{\sigma_t(p)}$ back to p . The resulting vector, shown in red, is $((\sigma_{-t})_* Y)_{\sigma_t(p)}$.

Constructing the Lie Derivative

Now we can use these ideas to help build a derivative. Suppose that we are given a vector field X on M . This generates a flow $\sigma_t : M \rightarrow M$, which is a map between manifolds, now with $N = M$. This means that we can use (2.11) to generate a push-forward map from $T_p(M)$ to $T_{\sigma_t(p)}(M)$. But this is exactly what we need if we want to compare tangent vectors at neighbouring points. The resulting differential operator is called the *Lie derivative* and is denoted \mathcal{L}_X .

It will turn out that we can use these ideas to differentiate many different kinds of objects. As a warm-up, let's first see how an analogous construction allows us to differentiate functions. Now the function

$$\mathcal{L}_X f = \lim_{t \rightarrow 0} \frac{f(\sigma_t(x)) - f(x)}{t} = \left. \frac{df(\sigma_t(x))}{dt} \right|_{t=0} = \left. \frac{\partial f}{\partial x^\mu} \frac{dx^\mu}{dt} \right|_{t=0}$$

But, using (2.9), we know that $dx^\mu/dt = X^\mu$. We then have

$$\mathcal{L}_X f = X^\mu(x) \frac{\partial f}{\partial x^\mu} = X(f) \tag{2.12}$$

In other words, acting on functions with the Lie derivative \mathcal{L}_X coincides with action of the vector field X .

Now let's look at the action of \mathcal{L}_X on a vector field Y . This is defined by

$$\mathcal{L}_X Y = \lim_{t \rightarrow 0} \frac{((\sigma_{-t})_* Y)_{\sigma_t(p)} - Y_p}{t}$$

Note the minus sign in σ_{-t} . This reflects that fact that vector fields are pushed, rather than pulled. The map σ_t takes us from the point p to the point $\sigma_t(p)$. But to push

a tangent vector $Y_{\sigma_t(p)} \in T_{\sigma_t(p)}(M)$ to a tangent vector in $T_p(M)$, where it can be compared to Y_p , we need to push with the inverse map $(\sigma_{-t})_*$. This is shown Figure 20.

Let's first calculate the action of \mathcal{L}_X on a coordinate basis $\partial_\mu = \partial/\partial x^\mu$. We have

$$\mathcal{L}_X \partial_\mu = \lim_{t \rightarrow 0} \frac{(\sigma_{-t})_* \partial_\mu - \partial_\mu}{t} \quad (2.13)$$

We have an expression for the push-forward of a tangent vector in (2.11), where the coordinates y^α on N should now be replaced by the infinitesimal change of coordinates induced by the flow σ_{-t} which, from (2.10) is $x^\mu(t) = x^\mu(0) - tX^\mu + \dots$. Note the minus sign, which comes from the fact that we have to map back to where we came from as shown in Figure 20. We have, for small t ,

$$(\sigma_{-t})_* \partial_\mu = \left(\delta_\mu^\nu - t \frac{\partial X^\nu}{\partial x^\mu} + \dots \right) \partial_\nu$$

Acting on a coordinate basis, we then have

$$\mathcal{L}_X \partial_\mu = - \frac{\partial X^\nu}{\partial x^\mu} \partial_\nu \quad (2.14)$$

To determine the action of \mathcal{L}_X on a general vector field Y , we use the fact that the Lie derivative obeys the usual properties that we expect of a derivative, including linearity, $\mathcal{L}_X(Y_1 + Y_2) = \mathcal{L}_X Y_1 + \mathcal{L}_X Y_2$ and Leibnizarity $\mathcal{L}_X(fY) = f\mathcal{L}_X Y + (\mathcal{L}_X f)Y$ for any function f , both of which follow from the definition. The action on a general vector field $Y = Y^\mu(x)\partial/\partial x^\mu$ can then be written as

$$\mathcal{L}_X(Y^\mu \partial_\mu) = (\mathcal{L}_X Y^\mu) \partial_\mu + Y^\mu (\mathcal{L}_X \partial_\mu)$$

where we've simply viewed the components $Y^\mu(x)$ as n functions. We can use (2.12) to determine $\mathcal{L}_X Y^\mu$ and we've computed $\mathcal{L}_X \partial_\mu$ in (2.14). We then have

$$\mathcal{L}_X(Y^\mu \partial_\mu) = X^\nu \frac{\partial Y^\mu}{\partial x^\nu} \partial_\mu - Y^\mu \frac{\partial X^\nu}{\partial x^\mu} \partial_\nu$$

But this is precisely the structure of the commutator. We learn that the Lie derivative acting on vector fields is given by

$$\mathcal{L}_X Y = [X, Y]$$

A corollary of this is

$$\mathcal{L}_X \mathcal{L}_Y Z - \mathcal{L}_Y \mathcal{L}_X Z = \mathcal{L}_{[X, Y]} Z \quad (2.15)$$

which follows from the Jacobi identity for commutators.

The Lie derivative is just one of several derivatives that we will meet in this course. As we introduce new objects, we will learn how to act with \mathcal{L}_X on them. But we will also see that we can endow different meanings to the idea of differentiation. In fact, the Lie derivative will take something of a back seat until Section 4.3 when we will see that it is what we need to understand symmetries.

2.3 Tensors

For any vector space V , the dual vector space V^* is the space of all linear maps from V to \mathbf{R} .

This is a standard mathematical construction, but even if you haven't seen it before it should resonate with something you know from quantum mechanics. There we have states in a Hilbert space with kets $|\psi\rangle \in \mathcal{H}$ and a dual Hilbert space with bras $\langle\phi| \in \mathcal{H}^*$. Any bra can be viewed as a map $\langle\phi| : \mathcal{H} \rightarrow \mathbf{R}$ defined by $\langle\phi|(|\psi\rangle) = \langle\phi|\psi\rangle$.

In general, suppose that we are given a basis $\{e_\mu, \mu = 1, \dots, n\}$ of V . Then we can introduce a dual basis $\{f^\mu, \mu = 1, \dots, n\}$ for V^* defined by

$$f^\nu(e_\mu) = \delta_\mu^\nu$$

A general vector in V can be written as $X = X^\mu e_\mu$ and $f^\nu(X) = X^\mu f^\nu(e_\mu) = X^\nu$. Given a basis, this construction provides an isomorphism between V and V^* given by $e_\mu \rightarrow f^\mu$. But the isomorphism is basis dependent. Pick a different basis, and you'll get a different map.

We can repeat the construction and consider $(V^*)^*$, which is the space of all linear maps from V^* to \mathbf{R} . But this space is naturally isomorphic to V , meaning that the isomorphism is independent of the choice of basis. To see this, suppose that $X \in V$ and $\omega \in V^*$. This means that $\omega(X) \in \mathbf{R}$. But we can equally as well view $X \in (V^*)^*$ and define $X(\omega) = \omega(X) \in \mathbf{R}$. In this sense, $(V^*)^* = V$.

2.3.1 Covectors and One-Forms

At each point $p \in M$, we have a vector space $T_p(M)$. The dual of this space, $T_p^*(M)$ is called the *cotangent space* at p , and an element of this space is called a *cotangent vector*, sometimes shortened to *covector*. Given a basis $\{e_\mu\}$ of $T_p(M)$, we can introduce the dual basis $\{f^\mu\}$ for $T_p^*(M)$ and expand any co-vector as $\omega = \omega_\mu f^\mu$.

We can also construct fields of cotangent vectors, by picking a member of $T_p^*(M)$ for each point p in a smooth manner. Such a *cotangent field* is better known as a *one-form*; they map vector fields to real numbers. The set of all one-forms on M is denoted $\Lambda^1(M)$.

There is a particularly simple way to construct a one-form. Take a function $f \in C^\infty(M)$ and define $df \in \Lambda^1(M)$ by

$$df(X) = X(f) \quad (2.16)$$

We can use this method to build a basis for $\Lambda^1(M)$. We introduce coordinates x^μ on M with the corresponding coordinate basis $e_\mu = \partial/\partial x^\mu$ of vector fields, which we often write in shorthand as $\partial/\partial x^\mu \equiv \partial_\mu$. We then simply take the functions $f = x^\mu$ which, from (2.16), gives

$$dx^\mu(\partial_\nu) = \partial_\nu(x^\mu) = \delta_\nu^\mu$$

This means that $f^\mu = dx^\mu$ provides a basis for $\Lambda^1(M)$, dual to the coordinate basis $\partial/\partial x^\mu$. In general, an arbitrary one-form $\omega \in \Lambda^1(M)$ can then be expanded as

$$\omega = \omega_\mu dx^\mu$$

In such a basis the one-form df takes the form

$$df = \frac{\partial f}{\partial x^\mu} dx^\mu \quad (2.17)$$

To see this, we simply need to evaluate

$$df(X) = \frac{\partial f}{\partial x^\mu} dx^\mu(X^\nu \partial_\nu) = X^\mu \frac{\partial f}{\partial x^\mu} = X(f)$$

which agrees with the expected answer (2.16).

As with vector fields, we can look at what happens if we change coordinates. Given two different charts, $\phi = (x^1, \dots, x^n)$ and $\tilde{\phi} = (\tilde{x}^1, \dots, \tilde{x}^n)$, we know that the basis for vector fields changes as (2.5),

$$\frac{\partial}{\partial \tilde{x}^\mu} = \frac{\partial x^\nu}{\partial \tilde{x}^\mu} \frac{\partial}{\partial x^\nu}$$

We should take the basis of one-forms to transform in the inverse manner,

$$d\tilde{x}^\mu = \frac{\partial \tilde{x}^\mu}{\partial x^\nu} dx^\nu \quad (2.18)$$

This then ensures that

$$d\tilde{x}^\mu \left(\frac{\partial}{\partial \tilde{x}^\nu} \right) = \frac{\partial \tilde{x}^\mu}{\partial x^\rho} dx^\rho \left(\frac{\partial x^\sigma}{\partial \tilde{x}^\nu} \frac{\partial}{\partial x^\sigma} \right) = \frac{\partial \tilde{x}^\mu}{\partial x^\rho} \frac{\partial x^\sigma}{\partial \tilde{x}^\nu} dx^\rho \left(\frac{\partial}{\partial x^\sigma} \right) = \frac{\partial \tilde{x}^\mu}{\partial x^\rho} \frac{\partial x^\sigma}{\partial \tilde{x}^\nu} \delta_\sigma^\rho$$

But this is just the multiplication of a matrix and its inverse,

$$\frac{\partial \tilde{x}^\mu}{\partial x^\rho} \frac{\partial x^\rho}{\partial \tilde{x}^\nu} = \delta_\nu^\mu$$

So we find that

$$d\tilde{x}^\mu \left(\frac{\partial}{\partial \tilde{x}^\nu} \right) = \delta_\nu^\mu$$

as it should. We can then expand a one-form ω in either of these two bases,

$$\omega = \omega_\mu dx^\mu = \tilde{\omega}_\mu d\tilde{x}^\mu \quad \text{with} \quad \tilde{\omega}_\mu = \frac{\partial x^\nu}{\partial \tilde{x}^\mu} \omega_\nu \quad (2.19)$$

In the annoying terminology that I can never remember, components of vectors that transform this way are said to be *covariant*. Note that, as with vector fields, the placement of the indices means that (2.18) and (2.19) are pretty much the only things that you can write down that make sense.

2.3.2 The Lie Derivative Revisited

In Section 2.2.4, we explained how to construct the Lie derivative, which differentiates a vector field in the direction of a second vector field X . This same idea can be adapted to one-forms.

Under a map $\varphi : M \rightarrow N$, we saw that a vector field X on M can be pushed forwards to a vector field $\varphi_* X$ on N . In contrast, one-forms go the other way: given a one-form ω on N , we can pull this back to a one-form $(\varphi^* \omega)$ on M , defined by

$$(\varphi^* \omega)(X) = \omega(\varphi_* X)$$

If we introduce coordinates x^μ on M and y^α on N then the components of the pull-back are given by

$$(\varphi^* \omega)_\mu = \omega_\alpha \frac{\partial y^\alpha}{\partial x^\mu} \quad (2.20)$$

We now define the Lie derivative \mathcal{L}_X acting on one-forms. Again, we use X to generate a flow $\sigma_t : M \rightarrow M$ which, using the pull-back, allows us to compare one-forms at different points. We will denote the cotangent vector $\omega(p)$ as ω_p . The Lie derivative of a one-form ω is then defined as

$$\mathcal{L}_X \omega = \lim_{t \rightarrow 0} \frac{(\sigma_t^* \omega)_p - \omega_p}{t} \quad (2.21)$$

Note that we pull-back with the map σ_t . This is to be contrasted with (2.13) where we pushed forward the tangent vector with the map σ_{-t} and, as we now show, this difference in minus sign manifests itself in the expression for the Lie derivative. The infinitesimal map σ_t acts on coordinates as $x^\mu(t) = x^\mu(0) + tX^\mu + \dots$ so, from (2.20), the pull-back of a basis vector dx^μ is

$$\sigma_t^* dx^\mu = \left(\delta_\nu^\mu + t \frac{\partial X^\mu}{\partial x^\nu} + \dots \right) dx^\nu$$

Acting on the coordinate basis, we then have

$$\mathcal{L}_X(dx^\mu) = \frac{\partial X^\mu}{\partial x^\nu} dx^\nu$$

which indeed differs by a minus sign from the corresponding result (2.14) for tangent vectors. Acting on a general one-form $\omega = \omega_\mu dx^\mu$, the Lie derivative is

$$\begin{aligned} \mathcal{L}_X \omega &= (\mathcal{L}_X \omega_\mu) dx^\mu + \omega_\nu \mathcal{L}_X(dx^\nu) \\ &= (X^\nu \partial_\nu \omega_\mu + \omega_\nu \partial_\mu X^\nu) dx^\mu \end{aligned} \tag{2.22}$$

We'll return to discuss one-forms (and other forms) more in Section 2.4.

2.3.3 Tensors and Tensor Fields

A *tensor of rank* (r, s) at a point $p \in M$ is defined to be a multi-linear map

$$T : \overbrace{T_p^*(M) \times \dots \times T_p^*(M)}^r \times \overbrace{T_p(M) \times \dots \times T_p(M)}^s \rightarrow \mathbf{R}$$

Such a tensor is said to have total rank $r + s$.

We've seen some examples already. A cotangent vector in $T_p^*(M)$ is a tensor of type $(0, 1)$, while a tangent vector in $T_p(M)$ is a tensor of type $(1, 0)$ (using the fact that $T_p(M) = T_p^{**}(M)$).

As before, we define a *tensor field* to be a smooth assignment of an (r, s) tensor to every point $p \in M$.

Given a basis $\{e_\mu\}$ for vector fields and a dual basis $\{f^\mu\}$ for one-forms, the components of the tensor are defined to be

$$T^{\mu_1 \dots \mu_r}_{\nu_1 \dots \nu_s} = T(f^{\mu_1}, \dots, f^{\mu_r}, e_{\nu_1}, \dots, e_{\nu_s})$$

Note that we deliberately write the string of lower indices after the upper indices. In some sense this is unnecessary, and we don't lose any information by writing $T^{\mu_1 \dots \mu_r}_{\nu_1 \dots \nu_s}$. Nonetheless, we'll see later that it's a useful habit to get into.

On a manifold of dimension n , there are n^{r+s} such components. For a tensor field, each of these is a function over M .

As an example, consider a rank $(2, 1)$ tensor. This takes two one-forms, say ω and η , together with a vector field X , and spits out a real number. In a given basis, this number is

$$T(\omega, \eta, X) = T(\omega_\mu f^\mu, \eta_\nu f^\nu, X^\rho e_\rho) = \omega_\mu \eta_\nu X^\rho T(f^\mu, f^\nu, e_\rho) = T^{\mu\nu}{}_\rho \omega_\mu \eta_\nu X^\rho$$

Every manifold comes equipped with a natural $(1, 1)$ tensor called δ . This takes a one-form ω and a vector field X and spits out the real number

$$\delta(\omega, X) = \omega(X) \quad \Rightarrow \quad \delta(f^\mu, e_\nu) = f^\mu(e_\nu) = \delta_\nu^\mu$$

which is simply the Kronecker delta.

As with vector fields and one-forms, we can ask how the components of a tensor transform. We will work more generally than before. Consider two bases for the vector fields, $\{e_\mu\}$ and $\{\tilde{e}_\mu\}$, not necessarily coordinate bases, related by

$$\tilde{e}_\nu = A^\mu{}_\nu e_\mu$$

for some invertible matrix A . The respective dual bases are $\{f^\mu\}$ and $\{\tilde{f}^\mu\}$ are then related by

$$\tilde{f}^\rho = B^\rho{}_\sigma f^\sigma$$

such that

$$\tilde{f}^\rho(\tilde{e}_\nu) = A^\mu{}_\nu B^\rho{}_\sigma f^\sigma(e_\mu) = A^\mu{}_\nu B^\rho{}_\mu = \delta_\nu^\rho \quad \Rightarrow \quad B^\rho{}_\mu = (A^{-1})^\rho{}_\mu$$

The lower components of a tensor then transform by multiplying by A , and the upper components by multiplying by $B = A^{-1}$. So, for example, a rank $(1, 2)$ tensor transforms as

$$\tilde{T}^\mu{}_{\rho\nu} = B^\mu{}_\sigma A^\tau{}_\rho A^\lambda{}_\nu T^\sigma{}_{\tau\lambda} \tag{2.23}$$

When we change between coordinate bases, we have

$$A^\mu{}_\nu = \frac{\partial x^\mu}{\partial \tilde{x}^\nu} \quad \text{and} \quad B^\mu{}_\nu = (A^{-1})^\mu{}_\nu = \frac{\partial \tilde{x}^\mu}{\partial x^\nu}$$

You can check that this coincides with our previous results (2.6) and (2.19).

Operations on Tensor Fields

There are a number of operations that we can do on tensor fields to generate further tensors.

First, we can add and subtract tensors fields, or multiply them by functions. This is the statement that the set of tensors at a point $p \in M$ forms a vector space.

Next, there is a way to multiply tensors together to give a tensor of a different type. Given a tensor S of rank (p, q) and a tensor T of rank (r, s) , we can form the *tensor product*, $S \otimes T$ which a new tensor of rank $(p + r, q + s)$, defined by

$$\begin{aligned} S \otimes T(\omega_1, \dots, \omega_p, \eta_1, \dots, \eta_r, X_1, \dots, X_q, Y_1, \dots, Y_s) \\ = S(\omega_1, \dots, \omega_p, X_1, \dots, X_q)T(\eta_1, \dots, \eta_r, Y_1, \dots, Y_s) \end{aligned}$$

In terms of components, this reads

$$(S \otimes T)^{\mu_1 \dots \mu_p \nu_1 \dots \nu_r}_{\rho_1 \dots \rho_q \sigma_1 \dots \sigma_s} = S^{\mu_1 \dots \mu_p}_{\rho_1 \dots \rho_q} T^{\nu_1 \dots \nu_r}_{\sigma_1 \dots \sigma_s} \quad (2.24)$$

Given an (r, s) tensor T , we can also construct a tensor of lower rank $(r - 1, s - 1)$ by contraction. To do this, simply replace one of $T_p^*(M)$ entries with a basis vector f^μ , and the corresponding $T_p(M)$ entry with the dual vector e_μ and then sum over $\mu = 1, \dots, n$. So, for example, given a rank $(2, 1)$ tensor T we can construct a rank $(1, 0)$ tensor S by

$$S(\omega) = T(\omega, f^\mu, e_\mu)$$

Alternatively, we could construct a (typically) different $(1, 0)$ tensor by contracting the other argument, $S'(\omega) = T(f^\mu, \omega, e_\mu)$. Written in terms of components, contraction simply means that we put an upper index equal to a lower index and sum over them,

$$S^\mu = T^{\mu\nu}{}_\nu \quad \text{and} \quad S'^\mu = T^{\nu\mu}{}_\nu$$

Our next operation is symmetrisation and anti-symmetrisation. For example, given a $(0, 2)$ tensor T we decompose it into two $(0, 2)$ tensors, in which the arguments are either symmetrised or anti-symmetrised,

$$\begin{aligned} S(X, Y) &= \frac{1}{2} \left(T(X, Y) + T(Y, X) \right) \\ A(X, Y) &= \frac{1}{2} \left(T(X, Y) - T(Y, X) \right) \end{aligned}$$

In index notation, this becomes

$$S_{\mu\nu} = \frac{1}{2}(T_{\mu\nu} + T_{\nu\mu}) \quad \text{and} \quad A_{\mu\nu} = \frac{1}{2}(T_{\mu\nu} - T_{\nu\mu})$$

which is just like taking the symmetric and anti-symmetric part of a matrix. We will work with these operations frequently enough to justify introducing some new notation. We define

$$T_{(\mu\nu)} = \frac{1}{2}(T_{\mu\nu} + T_{\nu\mu}) \quad \text{and} \quad T_{[\mu\nu]} = \frac{1}{2}(T_{\mu\nu} - T_{\nu\mu})$$

These operations generalise to other tensors. For example,

$$T^{(\mu\nu)\rho}{}_{\sigma} = \frac{1}{2}(T^{\mu\nu\rho}{}_{\sigma} + T^{\nu\mu\rho}{}_{\sigma})$$

We can also symmetrise or anti-symmetrise over multiple indices, provided that these indices are either all up or all down. If we (anti)-symmetrise over p objects, then we divide by $p!$, which is the number of possible permutations. This normalisation ensures that if we start with a tensor which is already, say, symmetric then further symmetrising doesn't affect it. In the case of anti-symmetry, we weight each term with the sign of the permutation. So, for example,

$$T^{\mu}{}_{(\nu\rho\sigma)} = \frac{1}{3!}(T^{\mu}{}_{\nu\rho\sigma} + T^{\mu}{}_{\rho\nu\sigma} + T^{\mu}{}_{\rho\sigma\nu} + T^{\mu}{}_{\sigma\rho\nu} + T^{\mu}{}_{\sigma\nu\rho} + T^{\mu}{}_{\nu\sigma\rho})$$

and

$$T^{\mu}{}_{[\nu\rho\sigma]} = \frac{1}{3!}(T^{\mu}{}_{\nu\rho\sigma} - T^{\mu}{}_{\rho\nu\sigma} + T^{\mu}{}_{\rho\sigma\nu} - T^{\mu}{}_{\sigma\rho\nu} + T^{\mu}{}_{\sigma\nu\rho} - T^{\mu}{}_{\nu\sigma\rho})$$

There will be times when, annoyingly, we will wish to symmetrise (or anti-symmetrise) over indices which are not adjacent. We introduce vertical bars to exclude certain indices from the symmetry procedure. So, for example,

$$T^{\mu}{}_{[\nu|\rho|\sigma]} = \frac{1}{2}(T^{\mu}{}_{\nu\rho\sigma} - T^{\mu}{}_{\sigma\rho\nu})$$

Finally, given a smooth tensor field T of any rank, we can always take the Lie derivative with respect to a vector field X . As we've seen previously, under a map $\varphi : M \rightarrow N$, vector fields are pushed forwards and one-forms are pulled-back. In general, this leaves a tensor of mixed type unsure where to go. However, if φ is a diffeomorphism then we also have $\varphi^{-1} : N \rightarrow M$ and this allows us to define the push-forward of a tensor T from M to N . This acts on one-forms $\omega \in \Lambda^1(N)$ and vector fields $X \in \mathfrak{X}(N)$ and is given by

$$(\varphi_*T)(\omega_1, \dots, \omega_r, X_1, \dots, X_s) = T(\varphi^*\omega_1, \dots, \varphi^*\omega_r, (\varphi_*^{-1}X_1), \dots, (\varphi_*^{-1}X_s))$$

Here $\varphi^*\omega$ are the pull-backs of ω from N to M , while $\varphi_*^{-1}X$ are the push-forwards of X from N to M .

The Lie derivative of a tensor T along X is then defined as

$$\mathcal{L}_X T = \lim_{t \rightarrow 0} \frac{((\sigma_{-t})_* T)_p - T_p}{t}$$

where σ_t is the flow generated by X . This coincides with our earlier definitions for vector fields in (2.13) and for one-forms in (2.21). (The difference in the σ_{-t} vs σ_t minus sign in (2.13) and (2.21) is now hiding in the inverse push-forward φ_*^{-1} that appears in the definition $\varphi_* T$.)

2.4 Differential Forms

Some tensors are more interesting than others. A particularly interesting class are totally anti-symmetric $(0, p)$ tensors fields. These are called p -forms. The set of all p -forms over a manifold M is denoted $\Lambda^p(M)$.

We've met some forms before. A 0-form is simply a function. Meanwhile, as we saw previously, a 1-form is another name for a covector. The anti-symmetry means that we can't have any form of degree $p > n = \dim(M)$. A p -form has $\binom{n}{p}$ different components. Forms in $\Lambda^n(M)$ are called *top forms*.

Given a p -form ω and a q -form η , we can take the tensor product (2.24) to construct a $(p+q)$ -tensor. If we anti-symmetrise this, we then get a $(p+q)$ -form. This construction is called the *wedge product*, and is defined by

$$(\omega \wedge \eta)_{\mu_1 \dots \mu_p \nu_1 \dots \nu_q} = \frac{(p+q)!}{p!q!} \omega_{[\mu_1 \dots \mu_p} \eta_{\nu_1 \dots \nu_q]}$$

where the $[\dots]$ in the subscript tells us to anti-symmetrise over all indices. For example, given $\omega, \eta \in \Lambda^1(M)$, we can construct a 2-form

$$(\omega \wedge \eta)_{\mu\nu} = \omega_\mu \eta_\nu - \omega_\nu \eta_\mu$$

For one forms, the anti-symmetry ensures that $\omega \wedge \omega = 0$. In general, if $\omega \in \Lambda^p(M)$ and $\eta \in \Lambda^q(M)$, then one can show that

$$\omega \wedge \eta = (-1)^{pq} \eta \wedge \omega$$

This means that $\omega \wedge \omega = 0$ for any form of odd degree. We can, however, wedge even degree forms with themselves. (Which you know already for 0-forms where the wedge product is just multiplication of functions.)

As a more specific example, consider $M = \mathbf{R}^3$ and $\omega = \omega_\mu dx^\mu$ and $\eta = \eta_\mu dx^\mu$. We then have

$$\begin{aligned}\omega \wedge \eta &= (\omega_1 dx^1 + \omega_2 dx^2 + \omega_3 dx^3) \wedge (\eta_1 dx^1 + \eta_2 dx^2 + \eta_3 dx^3) \\ &= (\omega_1 \eta_2 - \omega_2 \eta_1) dx^1 \wedge dx^2 + (\omega_2 \eta_3 - \omega_3 \eta_2) dx^2 \wedge dx^3 + (\omega_3 \eta_1 - \omega_1 \eta_3) dx^3 \wedge dx^1\end{aligned}$$

Notice that the components that arise are precisely those of the cross-product acting on vectors in \mathbf{R}^3 . This is no coincidence: what we usually think of as the cross-product between vectors is really a wedge product between forms. We'll have to wait to Section 3 to understand how to map from one to the other.

It can also be shown that the wedge product is associative, meaning

$$\omega \wedge (\eta \wedge \lambda) = (\omega \wedge \eta) \wedge \lambda$$

We can then drop the brackets in any such product.

Given a basis $\{f^\mu\}$ of $\Lambda^1(M)$, a basis of $\Lambda^p(M)$ can be constructed by wedge products $\{f^{\mu_1} \wedge \dots \wedge f^{\mu_p}\}$. We will usually work with the coordinate basis $\{dx^\mu\}$. This means that any p -form ω can be written locally as

$$\omega = \frac{1}{p!} \omega_{\mu_1 \dots \mu_p} dx^{\mu_1} \wedge \dots \wedge dx^{\mu_p} \quad (2.25)$$

Although locally any p -form can be written as (2.25), this may not be true globally. This, and related issues, will become of some interest in Section 2.4.3.

2.4.1 The Exterior Derivative

We learned in Section 2.3.1 how to construct a one-form df from a function f . In a coordinate basis, this one-form has components (2.17),

$$df = \frac{\partial f}{\partial x^\mu} dx^\mu$$

We can extend this definition to higher forms. The *exterior derivative* is a map

$$d : \Lambda^p(M) \rightarrow \Lambda^{p+1}(M)$$

In local coordinates (2.25), the exterior derivative acts as

$$(d\omega) = \frac{1}{p!} \frac{\partial \omega_{\mu_1 \dots \mu_p}}{\partial x^\nu} dx^\nu \wedge dx^{\mu_1} \wedge \dots \wedge dx^{\mu_p} \quad (2.26)$$

Equivalently we have

$$(d\omega)_{\mu_1 \dots \mu_{p+1}} = (p+1) \partial_{[\mu_1} \omega_{\mu_2 \dots \mu_{p+1}]} \quad (2.27)$$

Importantly, if we subsequently act with the exterior derivative again, we get

$$d(d\omega) = 0$$

because the derivatives are anti-symmetrised and hence vanish. This holds true for any p -form, a fact which is sometimes expressed as

$$d^2 = 0$$

It can be shown that the exterior derivative satisfies a number of further properties,

- $d(\omega \wedge \eta) = d\omega \wedge \eta + (-1)^p \omega \wedge d\eta$, where $\omega \in \Lambda^p(M)$.
- $d(\varphi^* \omega) = \varphi^*(d\omega)$ where φ^* is the pull-back associated to the map between manifolds, $\varphi : M \rightarrow N$
- Because the exterior derivative commutes with the pull-back, it also commutes with the Lie derivative. This ensures that we have $d(\mathcal{L}_X \omega) = \mathcal{L}_X(d\omega)$.

A p -form ω is said to be *closed* if $d\omega = 0$ everywhere. It is *exact* if $\omega = d\eta$ everywhere for some η . Because $d^2 = 0$, an exact form is necessarily closed. The question of when the converse is true is interesting: we'll discuss this more in Section 2.4.3.

Examples

Suppose that we have a one-form $\omega = \omega_\mu dx^\mu$, the exterior derivative gives a 2-form

$$(d\omega)_{\mu\nu} = \partial_\mu \omega_\nu - \partial_\nu \omega_\mu \quad \Rightarrow \quad d\omega = \frac{1}{2} (\partial_\mu \omega_\nu - \partial_\nu \omega_\mu) dx^\mu \wedge dx^\nu$$

As a specific example of this example, suppose that we take the one-form to live on \mathbf{R}^3 , with

$$\omega = \omega_1 dx^1 + \omega_2 dx^2 + \omega_3 dx^3$$

Since this is a field, each of the components ω_μ is a function of x^1 , x^2 and x^3 . The exterior derivative is given by

$$\begin{aligned} d\omega &= \partial_2 \omega_1 dx^2 \wedge dx^1 + \partial_3 \omega_1 dx^3 \wedge dx^1 + \partial_1 \omega_2 dx^1 \wedge dx^2 \\ &\quad + \partial_3 \omega_2 dx^3 \wedge dx^2 + \partial_1 \omega_3 dx^1 \wedge dx^3 + \partial_2 \omega_3 dx^2 \wedge dx^3 \\ &= (\partial_1 \omega_2 - \partial_2 \omega_1) dx^1 \wedge dx^2 + (\partial_2 \omega_3 - \partial_3 \omega_2) dx^2 \wedge dx^3 + (\partial_3 \omega_1 - \partial_1 \omega_3) dx^3 \wedge dx^1 \end{aligned} \quad (2.28)$$

Notice that there's no term like $\partial_1 \omega_1$ because this would come with a $dx^1 \wedge dx^1 = 0$.

In the olden days (before this course), we used to write vector fields in \mathbf{R}^3 as $\boldsymbol{\omega} = (\omega^1, \omega^2, \omega^3)$ and compute the curl $\nabla \times \boldsymbol{\omega}$. But the components of the curl are precisely the components that appear in $d\omega$. In fact, our “vector” $\boldsymbol{\omega}$ was really a one-form and the curl turned it into a two-form. It’s a happy fact that in \mathbf{R}^3 , vectors, one-forms and two-forms all have three components, which allowed us to conflate them in our earlier courses. (In fact, there is a natural map between them that we will meet in Section 3.)

Suppose instead that we start with a 2-form B in \mathbf{R}^3 , which we write as

$$B = B_1 dx^2 \wedge dx^3 + B_2 dx^3 \wedge dx^1 + B_3 dx^1 \wedge dx^2$$

Taking the exterior derivative now gives

$$\begin{aligned} dB &= \partial_1 B_1 dx^1 \wedge dx^2 \wedge dx^3 + \partial_2 B_2 dx^2 \wedge dx^3 \wedge dx^1 + \partial_3 B_3 dx^3 \wedge dx^1 \wedge dx^2 \\ &= (\partial_1 B_1 + \partial_2 B_2 + \partial_3 B_3) dx^1 \wedge dx^2 \wedge dx^3 \end{aligned} \quad (2.29)$$

This time there is just a single component, but again it’s something familiar. Had we written the original three components of the two-form in old school vector notation $\mathbf{B} = (B_1, B_2, B_3)$, then the single component of dB is what we previously called $\nabla \cdot \mathbf{B}$.

The Lie Derivative Yet Again

There is yet another operation that we can construct on p -forms. Given a vector field $X \in \mathfrak{X}(M)$, we can construct the *interior product*, a map $\iota_X : \Lambda^p(M) \rightarrow \Lambda^{p-1}(M)$. If $\omega \in \Lambda^p(M)$, we define a $\iota_X \omega \in \Lambda^{p-1}(M)$ by

$$\iota_X \omega(Y_1, \dots, Y_{p-1}) = \omega(X, Y_1, \dots, Y_{p-1}) \quad (2.30)$$

In other words, we just put X in the first argument of ω . Acting on functions f , we simply define $\iota_X f = 0$.

The anti-symmetry of forms means that $\iota_X \iota_Y = -\iota_Y \iota_X$. Moreover, you can check that

$$\iota_X(\omega \wedge \eta) = \iota_X \omega \wedge \eta + (-1)^p \omega \wedge \iota_X \eta$$

where $\omega \in \Lambda^p(M)$.

Consider a 1-form ω . There are two different ways to act with ι_X and d to give us back a one-form. These are

$$\iota_X d\omega = \iota_X \frac{1}{2}(\partial_\mu \omega_\nu - \partial_\nu \omega_\mu) dx^\mu \wedge dx^\nu = X^\mu \partial_\mu \omega_\nu dx^\nu - X^\nu \partial_\nu \omega_\mu dx^\mu$$

and

$$d\iota_X\omega = d(\omega_\mu X^\mu) = X^\mu \partial_\nu \omega_\mu dx^\nu + \omega_\mu \partial_\nu X^\mu dx^\nu$$

Adding the two together gives

$$(d\iota_X + \iota_X d)\omega = (X^\mu \partial_\mu \omega_\nu + \omega_\mu \partial_\nu X^\mu) dx^\nu$$

But this is exactly the same expression we saw previously when computing the Lie derivative (2.22) of a one-form. We learn that

$$\mathcal{L}_X\omega = (d\iota_X + \iota_X d)\omega \tag{2.31}$$

This expression is sometimes referred to as *Cartan's magic formula*. A similar calculation shows that (2.31) holds for any p -form ω .

2.4.2 Forms You Know and Love

There are a number of examples of differential forms that you've met already, but likely never called them by name.

The Electromagnetic Field

The electromagnetic gauge field $A_\mu = (\phi, \mathbf{A})$ should really be thought of as the components of a one-form on spacetime \mathbf{R}^4 . (Here I've set $c = 1$.) We write

$$A = A_\mu(x) dx^\mu$$

Taking the exterior derivative yields a 2-form $F = dA$, given by

$$F = \frac{1}{2} F_{\mu\nu} dx^\mu \wedge dx^\nu = \frac{1}{2} (\partial_\mu A_\nu - \partial_\nu A_\mu) dx^\mu \wedge dx^\nu$$

But this is precisely the field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ that we met in our lectures on [Electromagnetism](#). The components are the electric and magnetic fields, arranged as

$$F_{\mu\nu} = \begin{pmatrix} 0 & -E_1 & -E_2 & -E_3 \\ E_1 & 0 & B_3 & -B_2 \\ E_2 & -B_3 & 0 & B_1 \\ E_3 & B_2 & -B_1 & 0 \end{pmatrix} \tag{2.32}$$

By construction, we also have $dF = d^2A = 0$. In this context, this is sometimes called the *Bianchi identity*; it yields two of the four Maxwell equations. In old school vector calculus notation, these are $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{E} + \partial\mathbf{B}/\partial t = 0$. We need a little more structure to get the other two as we will see later in this chapter.

The gauge field A is not unique. Given any function α , we can always shift it by a gauge transformation

$$A \rightarrow A + d\alpha \quad \Rightarrow \quad A_\mu \rightarrow A_\mu + \partial_\mu \alpha$$

This leaves the field strength invariant because $F \rightarrow F + d(d\alpha) = F$.

Phase Space and Hamilton's Equations

In classical mechanics, the phase space is a manifold M parameterised by coordinates (q^i, p_j) where q^i are the positions of particles and p_j the momenta. Recall from our lectures on [Classical Dynamics](#) that the Hamiltonian $H(q, p)$ is a function on M , and Hamilton's equations are

$$\dot{q}^i = \frac{\partial H}{\partial p_i} \quad \text{and} \quad \dot{p}_i = -\frac{\partial H}{\partial q^i} \quad (2.33)$$

Phase space also comes with a structure called a *Poisson bracket*, defined on a pair of functions f and g as

$$\{f, g\} = \frac{\partial f}{\partial q^j} \frac{\partial g}{\partial p_j} - \frac{\partial f}{\partial p_j} \frac{\partial g}{\partial q^j}$$

Then the time evolution of any function f can be written as

$$\dot{f} = \{f, H\}$$

which reproduces Hamilton's equations if we take $f = q^i$ or $f = p_i$.

Underlying this story is the mathematical structure of forms. The key idea is that we have a manifold M and a function H on M . We want a machinery that turns the function H into a vector field X_H . Particles then follow trajectories in phase space that are integral curves generated by X_H .

To achieve this, we introduce a *symplectic two-form* ω on an even-dimensional manifold M . This two form must be closed, $d\omega = 0$, and non-degenerate, which means that the top form $\omega \wedge \dots \wedge \omega \neq 0$. We'll see why we need these requirements as we go along. A manifold M equipped with a symplectic two-form is called a *symplectic manifold*.

Any 2-form provides a map $\omega : T_p(M) \rightarrow T_p^*(M)$. Given a vector field $X \in \mathfrak{X}(M)$, we can simply take the interior product with ω to get a one-form $\iota_X \omega$. However, we want to go the other way: given a function H , we can always construct a one-form dH , and we'd like to exchange this for a vector field. We can do this if the map

$\omega : T_p(M) \rightarrow T_p^*(M)$ is actually an isomorphism, so the inverse exists. This turns out to be true provided that ω is non-degenerate. In this case, we can define the vector field X_H by solving the equation

$$\iota_{X_H}\omega = -dH \quad (2.34)$$

If we introduce coordinates x^μ on the manifold, then the component form of this equation is

$$X_H^\mu \omega_{\mu\nu} = -\partial_\nu H$$

We denote the inverse as $\omega^{\mu\nu} = -\omega^{\nu\mu}$ such that $\omega^{\mu\nu}\omega_{\nu\rho} = \delta_\rho^\mu$. The components of the vector field are then

$$X_H^\mu = -\omega^{\nu\mu}\partial_\nu H = \omega^{\mu\nu}\partial_\nu H$$

The integral curves generated by X_H obey the differential equation (2.9)

$$\frac{dx^\mu}{dt} = X_H^\mu = \omega^{\mu\nu}\partial_\nu H$$

These are the general form of Hamilton's equations. They reduce to our earlier form (2.33) if we write $x^\mu = (q^i, p_j)$ and choose the symplectic form to have block diagonal form $\omega^{\mu\nu} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$.

To define the Poisson structure, we first note that we can repeat the map (2.34) to turn any function f into a vector field X_f obeying $\iota_{X_f}\omega = -df$. But we can then feed these vector fields back into the original 2-form ω . This gives us a Poisson bracket,

$$\{f, g\} = \omega(X_g, X_f) = -\omega(X_f, X_g)$$

Or, in components,

$$\{f, g\} = \omega^{\mu\nu}\partial_\mu f \partial_\nu g$$

There are many other ways to write this Poisson bracket structure in invariant form. For example, backtracking through various definitions we find

$$\{f, g\} = -\iota_{X_f}\omega(X_g) = df(X_g) = X_g(f)$$

The equation of motion in Poisson bracket structure is then

$$\dot{f} = \{f, H\} = X_H(f) = \mathcal{L}_{X_H}f$$

which tells us that the Lie derivative along X_H generates time evolution.

We haven't yet explained why the symplectic two-form must be closed, $d\omega = 0$. You can check that this is needed so that the Poisson bracket obeys the Jacobi identity. Alternatively, it ensures that the symplectic form itself is invariant under Hamiltonian flow, in the sense that $\mathcal{L}_{X_H}\omega = 0$. To see this, we use (2.31)

$$\mathcal{L}_{X_H}\omega = (d\iota_{X_H} + \iota_{X_H}d)\omega = \iota_{X_H}d\omega$$

The second equality follows from the fact that $d\iota_{X_H}\omega = -d(dH) = 0$. If we insist that $d\omega = 0$ then we find $\mathcal{L}_{X_H}\omega = 0$ as promised.

Thermodynamics

The state space of a thermodynamic system is a manifold M . For the ideal gas, this is a two-dimensional manifold with coordinates provided by, for example, the pressure p and volume V . More complicated systems can have a higher dimensional state space, but the dimension is always even since, like in classical mechanics, thermodynamical variables come in conjugate pairs.

When we first learn the [laws of thermodynamics](#), we have to make up some strange new notation, \vec{d} , which then never rears its head again. For example, the first law of thermodynamics is written as

$$dE = \vec{d}Q + \vec{d}W$$

Here dE is the infinitesimal change of energy in the system. The first law of thermodynamics, as written above, states that this decomposes into the heat flowing into the system $\vec{d}Q$ and the work done on the system $\vec{d}W$.

Why the stupid notation? Well, the energy $E(p, q)$ is a function over the state space M and this means that we can write the change of the energy $dE = \frac{\partial E}{\partial p}dp + \frac{\partial E}{\partial V}dV$. But there is no such function $Q(p, q)$ or $W(p, q)$ and, correspondingly, $\vec{d}Q$ and $\vec{d}W$ are not exact differentials. Indeed, we have $\vec{d}W = -pdV$ and later, after we introduce the second law, we learn that $\vec{d}Q = TdS$, with T the temperature and S the entropy, both of which are functions over M .

This is much more natural in the language of forms. All of the terms in the first law are one-forms. But the transfer of heat, $\vec{d}Q$ and the work $\vec{d}W$ are not exact one-forms and so can't be written as $d(\text{something})$. In contrast, dE is an exact one-form. That's what the \vec{d} notation is really telling us: it's the way of denoting non-exact one-forms before we had a notion of differential geometry.

The real purpose of the first law of thermodynamics is to define the energy functional E . The 18th century version of the statement is something like: “the amount of work required to change an isolated system is independent of how the work is performed”. A more modern rendering would be: “the sum of the work and heat is an exact one-form”.

2.4.3 A Sniff of de Rham Cohomology

The exterior derivative is a map which squares to zero, $d^2 = 0$. It turns out that one can have a lot of fun with such maps. We will now explore a little bit of this fun.

First a repeat of definitions we met already: a p -form is *closed* if $d\omega = 0$ everywhere. A p -form is *exact* if $\omega = d\eta$ everywhere for some η . Because $d^2 = 0$, exact implies closed. However, the converse is not necessarily true. It turns out that the way in which closed forms fail to be exact captures interesting facts about the topology of the underlying manifold.

We’ve met this kind of question before. In electromagnetism, we have a magnetic field \mathbf{B} which obeys $\nabla \cdot \mathbf{B} = 0$. We then argue that this means we can write the magnetic field as $\mathbf{B} = \nabla \times \mathbf{A}$. This is more properly expressed the language of forms. We we saw in the previous section, the magnetic field is really a 2-form

$$B = B_1 dx^2 \wedge dx^3 + B_2 dx^3 \wedge dx^1 + B_3 dx^1 \wedge dx^2$$

We computed the exterior derivative in (2.29); it is

$$dB = (\partial_1 B_1 + \partial_2 B_2 + \partial_3 B_3) dx^1 \wedge dx^2 \wedge dx^3$$

We see that the Maxwell equation $\nabla \cdot \mathbf{B} = 0$ is really the statement that B is a closed two-form, obeying $dB = 0$. We also saw in (2.28) if we write $B = dA$ for some one-form A , then the components are given by $\mathbf{B} = \nabla \times \mathbf{A}$. Clearly writing $B = dA$ ensures that $dB = 0$. But when is the converse true? We have the following statement (which we leave unproven)

Theorem (The Poincaré Lemma): On $M = \mathbf{R}^n$, closed implies exact.

Since we’ve spent a lot of time mapping manifolds to \mathbf{R}^n , this also has consequence for a general manifold M . It means that if ω is a closed p -form, then in any neighbourhood $\mathcal{O} \subset M$ it is always possible to find a $\eta \in \Lambda^{p-1}(M)$ such that $\omega = d\eta$ on \mathcal{O} . The catch is that it may not be possible to find such an η everywhere on the manifold.

An Example

Consider the one-dimensional manifold $M = \mathbf{R}$. We can take a one-form $\omega = f(x)dx$. This is always closed because it is a top form. It is also exact. We introduce the function

$$g(x) = \int_0^x dx' f(x')$$

Then $\omega = dg$.

Now consider the topologically more interesting one-dimensional manifold \mathbf{S}^1 , which we can view as the phase $e^{i\theta} \in \mathbf{C}$. We can introduce the form $\omega = d\theta$ on \mathbf{S}^1 . The way its written makes it look like its an exact form, but this is an illusion because, as we stressed in Section 2.1, θ is not a good coordinate everywhere on \mathbf{S}^1 because it's not single valued. Indeed, it's simple to see that there is no single-valued function $g(\theta)$ on \mathbf{S}^1 such that $\omega = dg$. So on \mathbf{S}^1 , we can construct a form which, locally, can be written as $d\theta$ but globally cannot be written as $d(\text{something})$. So we have a form that is closed but not exact.

Another Example

On $M = \mathbf{R}^2$, the Poincaré lemma ensures that all closed forms are exact. However, things change if we remove a single point and consider $\mathbf{R}^2 - \{0, 0\}$. Consider the one-form defined by

$$\omega = -\frac{y}{x^2 + y^2}dx + \frac{x}{x^2 + y^2}dy$$

This is not a smooth one-form on \mathbf{R}^2 because of the divergence at the origin. But removing that point means that ω becomes acceptable. We can check that ω is closed,

$$d\omega = -\frac{\partial}{\partial y} \left(\frac{y}{x^2 + y^2} \right) dy \wedge dx + \frac{\partial}{\partial x} \left(\frac{x}{x^2 + y^2} \right) dx \wedge dy = 0$$

where the $= 0$ follows from a little bit of algebra. ω is exact if we can find a function f , defined everywhere on $\mathbf{R}^2 - \{0, 0\}$ such that $\omega = df$, which means

$$\omega = \frac{\partial f}{\partial x}dx + \frac{\partial f}{\partial y}dy \quad \Rightarrow \quad \frac{\partial f}{\partial x} = -\frac{y}{x^2 + y^2} \quad \text{and} \quad \frac{\partial f}{\partial y} = \frac{x}{x^2 + y^2}$$

We can certainly integrate these equations; the result is

$$f(x, y) = \tan^{-1} \left(\frac{y}{x} \right) + \text{constant}$$

But this is not a smooth function everywhere on $\mathbf{R}^2 - \{0, 0\}$. This means that we can't, in fact, write $\omega = df$ for a well defined function on $\mathbf{R}^2 - \{0, 0\}$. We learn that removing a point makes a big difference: now closed no longer implies exact.

There is a similar story for \mathbf{R}^3 . Indeed, this is how magnetic monopoles sneak back into physics, despite being forbidden by the Maxwell equation $\nabla \cdot \mathbf{B} = 0$. You can learn more about this in the lectures on [Gauge Theory](#).

Betti Numbers

We denote the set of all closed p -forms on a manifold M as $Z^p(M)$. Equivalently, $Z^p(M)$ is the kernel of the map $d : \Lambda^p(M) \rightarrow \Lambda^{p+1}(M)$.

We denote the set of all exact p -forms on a manifold M as $B^p(M)$. Equivalently, $B^p(M)$ is the range of $d : \Lambda^{p-1}(M) \rightarrow \Lambda^p(M)$.

The p^{th} *de Rham cohomology group* is defined to be

$$H^p(M) = Z^p(M)/B^p(M)$$

The quotient here is an equivalence class. Two closed forms $\omega, \omega' \in Z^p(M)$ are said to be equivalent if $\omega = \omega' + \eta$ for some $\eta \in B^p(M)$. We say that ω and ω' sit in the same equivalence class $[\omega]$. The cohomology group $H^p(M)$ is the set of equivalence classes; in other words, it consists of closed forms mod exact forms.

The *Betti numbers* B_p of a manifold M are defined as

$$B_p = \dim H^p(M)$$

It turns out that these are always finite. The Betti number $B_0 = 1$ for any connected manifold. This can be traced to the existence of constant functions which are clearly closed but, because there are no $p = -1$ forms, are not exact. The higher Betti numbers are non-zero only if the manifold has some interesting topology. Finally, the *Euler character* is defined as the alternating sum of Betti numbers,

$$\chi(M) = \sum_p (-1)^p B_p \tag{2.35}$$

Here are some simple examples. We've already seen that the circle \mathbf{S}^1 has a closed, non-exact one-form. This means that $B_1 = 1$ and $\chi = 0$. The sphere \mathbf{S}^n has only $B_n = 1$ and $\chi = 1 + (-1)^n$. The torus \mathbf{T}^n has $B_p = \binom{n}{p}$ and $\chi = 0$.

2.4.4 Integration

We have learned how to differentiate on manifolds by using a vector field X . Now it is time to learn how to integrate. It turns out that the things that we integrate on manifolds are forms.

Integrating over Manifolds

To start, we need to orient ourselves. A *volume form*, or *orientation* on a manifold of dimension $\dim(M) = n$ is a nowhere-vanishing top form v . Any top form has just a single component and can be locally written as

$$v = v(x) dx^1 \wedge \dots \wedge dx^n$$

where we require $v(x) \neq 0$. If such a top form exists everywhere on the manifold, then M is said to be *orientable*.

The orientation is called right-handed if $v(x) > 0$ everywhere, and left-handed if $v(x) < 0$ everywhere. Given one volume form v , we can always construct another by multiplying by a function, giving $\tilde{v} = fv$ where $f(x) > 0$ everywhere or $f(x) < 0$ everywhere.

It's not enough to just write down a volume form with $v(x) \neq 0$ locally. We must also ensure that we can patch these volume forms together over the manifold, without the handedness changing. Suppose that we have two sets of coordinates, x^μ and \tilde{x}^μ that overlap on some region. In the new coordinates, the volume form is given by

$$v = v(x) \frac{\partial x^1}{\partial \tilde{x}^{\mu_1}} d\tilde{x}^{\mu_1} \wedge \dots \wedge \frac{\partial x^n}{\partial \tilde{x}^{\mu_n}} d\tilde{x}^{\mu_n} = v(x) \det \left(\frac{\partial x^\mu}{\partial \tilde{x}^\nu} \right) d\tilde{x}^1 \wedge \dots \wedge d\tilde{x}^n$$

which has the same orientation provided

$$\det \left(\frac{\partial x^\mu}{\partial \tilde{x}^\nu} \right) > 0 \tag{2.36}$$

Non-orientable manifolds cannot be covered by overlapping charts such that (2.36) holds. Examples include the Möbius strip and real projective space \mathbf{RP}^n for n even. (In contrast \mathbf{RP}^n is orientable for n odd, and \mathbf{CP}^n is orientable for all n .) In these lectures, we deal only with orientable manifolds.

Given a volume form v on M , we can integrate any function $f : M \rightarrow \mathbf{R}$ over the manifold. In a chart $\phi : \mathcal{O} \rightarrow U$, with coordinates x^μ , we have

$$\int_{\mathcal{O}} fv = \int_U dx_1 \dots dx_n f(x)v(x)$$

On the right-hand-side, we're just doing normal integration over some part of \mathbf{R}^n . The volume form is playing the role of a measure, telling us how much to weight various parts of the integral. To integrate over the entire manifold, we divide the manifold up into different regions, each covered by a single chart. We then perform the integral over each region and sum the results.

Integrating over Submanifolds

We don't have to integrate over the full manifold M . We can integrate over some lower dimensional submanifold.

A manifold Σ with dimension $k < n$ is a *submanifold* of M if we can find a map $\phi : \Sigma \rightarrow M$ which is one-to-one (which ensures that Σ doesn't intersect itself in M) and $\phi_* : T_p(\Sigma) \rightarrow T_{\phi(p)}(M)$ is one-to-one.

We can then integrate a k -form ω on M over a k -dimensional submanifold Σ . We do this by pulling back the k -form to Σ and writing

$$\int_{\phi(\Sigma)} \omega = \int_{\Sigma} \phi^* \omega$$

For example, suppose that we have a one-form A living over M . If C is a one-dimensional manifold, then we can introduce a map $\sigma : C \rightarrow M$ which defines a non-intersecting, one-dimensional curve $\sigma(C)$ which is a submanifold of M . We can then pull-back A onto this curve and integrate to get

$$\int_{\sigma(C)} A = \int_C \sigma^* A$$

This probably looks more familiar in coordinates. If the curve traces out a path $x^\mu(\tau)$ in M , we have

$$\int_C \sigma^* A = \int d\tau A_\mu(x) \frac{dx^\mu}{d\tau}$$

But this is precisely the way the worldline of a particle couples to the electromagnetic field, as we previously saw in (1.20).

2.4.5 Stokes' Theorem

Until now, we have considered only smooth manifolds. There is a slight generalisation that will be useful. We define a *manifold with boundary* in the same way as a manifold, except the charts map $\phi : \mathcal{O} \rightarrow U$ where U is an open subset of $\mathbf{R}^{n+} = \{\{x^1, \dots, x^n\} \text{ such that } x^n \geq 0\}$. The boundary has co-dimension 1 and is denoted ∂M : it is the submanifold with coordinates $x^n = 0$.

Consider a manifold M with boundary ∂M . If the dimension of the manifold is $\dim(M) = n$ then for any $(n - 1)$ -form ω , we have the following simple result

$$\int_M d\omega = \int_{\partial M} \omega \quad (2.37)$$

This is *Stokes' theorem*.

We do not prove Stokes' theorem here. The proof is fairly tedious, and does not differ greatly from the proofs of other things that you've called Stokes' theorem (or Gauss' divergence theorem) in the past. (See, for example, the lectures on [Vector Calculus](#).) However, the wonderful thing about (2.37) is the way in which it unifies many results in a single elegant formula. To see this, we simply need to look at a few examples.

The Mother of all Integral Theorems

First, consider $n = 1$ with M the interval I . We introduce coordinates $x \in [a, b]$ on the interval. The 0-form $\omega = \omega(x)$ is simply a function and $d\omega = (d\omega/dx)dx$. In this case, the two sides of Stokes' theorem can be evaluated to give

$$\int_M d\omega = \int_a^b \frac{d\omega}{dx} dx \quad \text{and} \quad \int_{\partial M} \omega = \omega(b) - \omega(a)$$

Equating the two, we see that Stokes' theorem is simply a restatement of the fundamental theorem of calculus.

Next, we take $M \subset \mathbf{R}^2$ to be a manifold with boundary. We introduce a one-form with coordinates

$$\omega = \omega_1 dx^1 + \omega_2 dx^2 \quad \Rightarrow \quad d\omega = \left(\frac{\partial \omega_2}{\partial x^1} - \frac{\partial \omega_1}{\partial x^2} \right) dx^1 \wedge dx^2$$

In this case, the ingredients in Stokes' theorem are

$$\int_M d\omega = \int_M \left(\frac{\partial \omega_2}{\partial x^1} - \frac{\partial \omega_1}{\partial x^2} \right) dx^1 dx^2 \quad \text{and} \quad \int_{\partial M} \omega = \int_{\partial M} \omega_1 dx^1 + \omega_2 dx^2$$

Equating the two gives the result usually referred to as Green's theorem in the plane.

Finally, consider $M \subset \mathbf{R}^3$ to be a manifold with boundary, with a 2-form

$$\omega = \omega_1 dx^2 \wedge dx^3 + \omega_2 dx^3 \wedge dx^1 + \omega_3 dx^1 \wedge dx^2$$

The right-hand-side of Stokes theorem is

$$\int_{\partial M} \omega_1 dx^2 dx^3 + \omega_2 dx^3 dx^1 + \omega_3 dx^1 dx^2$$

Meanwhile, we computed the exterior derivative of a 2-form in (2.29). The left-hand-side of Stokes' theorem then gives

$$\int_M d\omega = \int_M (\partial_1\omega_1 + \partial_2\omega_2 + \partial_3\omega_3) dx^1 dx^2 dx^3$$

This time, equating the two gives us Gauss' divergence theorem.

We see that Stokes' theorem, as written in (2.37), is the mother of all integral theorems, packaging many famous results in a single formula. We'll revisit this in Section 3.2.4 where we relate Stokes' theorem to a more explicit form of the divergence theorem.